

Quality-Diversity Generative Sampling for Learning with Synthetic Data

Allen Chang¹, Matthew C. Fontaine¹, Serena Booth², Maja J. Matarić¹, Stefanos Nikolaidis¹

¹ University of Southern California, Los Angeles, USA

² Massachusetts Institute of Technology, Cambridge, USA

{changall, mfontain, mataric, nikolaid}@usc.edu, sbooth@mit.edu

Abstract

Generative models can serve as surrogates for some real data sources by creating synthetic training datasets, but in doing so they may transfer biases to downstream tasks. We focus on protecting quality and diversity when generating synthetic training datasets. We propose quality-diversity generative sampling (QDGS), a framework for sampling data uniformly across a user-defined measure space, despite the data coming from a biased generator. QDGS is a model-agnostic framework that uses prompt guidance to optimize a quality objective across measures of diversity for synthetically generated data, without fine-tuning the generative model. Using balanced synthetic datasets generated by QDGS, we first debias classifiers trained on color-biased shape datasets as a proof-of-concept. By applying QDGS to facial data synthesis, we prompt for desired semantic concepts, such as skin tone and age, to create an intersectional dataset with a combined blend of visual features. Leveraging this balanced data for training classifiers improves fairness while maintaining accuracy on facial recognition benchmarks. Code available at: <https://github.com/Cylumn/qd-generative-sampling>.

1 Introduction

Generative models can produce large amounts of synthetic data at low cost, making them viable surrogates for some real data sources in model training (Azizi et al. 2023; Tian et al. 2023). However, generative models suffer from representational biases (Cho, Zala, and Bansal 2023; Friedrich et al. 2023; Bau et al. 2019), endangering a transfer of bias to downstream tasks (Jain et al. 2022). We propose a sampling framework that protects quality and diversity when generating synthetic datasets, and we show that these datasets can suppress performance decline on minority groups (Fig. 1).

Balancing synthetic datasets is not a trivial task, since generative models harbor biases that relate to different semantic concepts. Some biases result from training deep generative models to replicate Internet data, where there are common issues of sampling selection bias. Internet data also contain social biases that result from disparities in language use among social groups, eliciting stereotypes and prejudices (Otterbacher 2018). Biases are intensified by selecting training approaches that amplify class and feature imbalances, such as some data augmentation and regularization

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

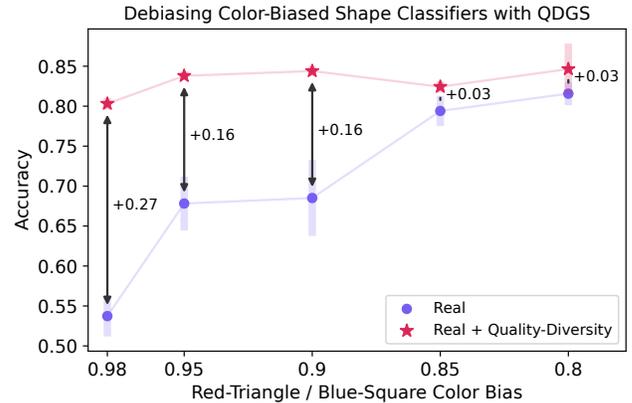


Figure 1: Fine-tuning color-biased shape classifiers with balanced data generated from QDGS results in higher accuracy across minority shape-color combinations. Standard error bars across 5 trials are shaded in.

methods (Balestriero, Bottou, and LeCun 2022). Thus, most existing generative models suffer from multiple biases that harm the creation of fair synthetic datasets. When creating these datasets for training downstream models, it is helpful to simultaneously consider multiple of these biases.

However, prior methods for sampling debiased synthetic data tend to focus on removing a single bias at a time, through single-factor diversity optimization (McDuff et al. 2019; Jain, Memon, and Togelius 2023). By increasing diversity along one dimension at a time, sampling methods miss the opportunity to generate intersectional datasets.

We leverage the insight that a generative model’s high dimensionality latent space can often combine underrepresented semantic concepts to create synthetic datasets with a rich blend of intersectional data. However, the same latent spaces typically collapse on modes that disproportionately represent the dominant classes or data attributes. To address these biases, we propose exploring the latent space to identify and generate underrepresented attribute combinations. We call this approach *quality-diversity generative sampling (QDGS)*, a framework using quality-diversity optimization with prompt guidance for producing synthetic datasets with a balanced spread of desired features.

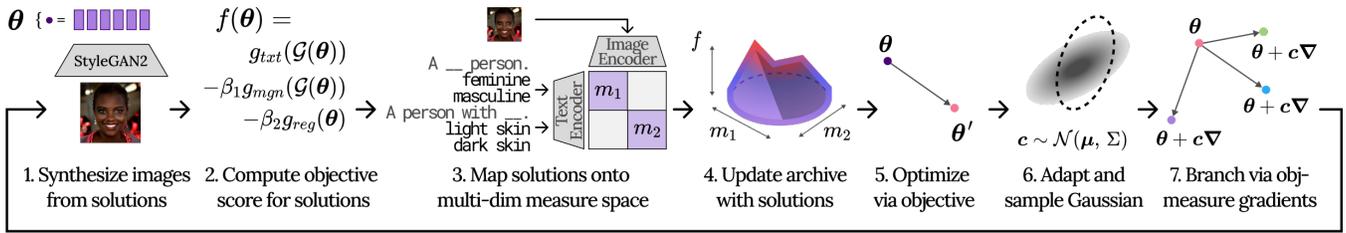


Figure 2: An implementation of the data generation pipeline used in QDGS. Steps 1-4 describe how solutions from QDGS are mapped to the objective-measure archive, and steps 5-7 describe how solutions are updated.

To better understand how to improve parity across different semantic concepts, we consider the domain of facial images. Randomly sampling from StyleGAN2 (Karras et al. 2020) reproduces a skin tone imbalance ($\sim 7:1$ light¹ to dark skin tone ratio) and an age imbalance ($\sim 2:1$ young to old ratio) from its training dataset Flickr-Faces-HQ (FFHQ) (Karras, Laine, and Aila 2019). Synthetic facial datasets typically contain other demographic and visual disparities (Table 1), which can create multiple dimensions of bias.

By employing user-defined language prompts to describe diversity measures (e.g., “dark skin tone,” “smiling,” “dim lighting,” etc.) and an objective (e.g., “detailed image”), QDGS enables a greater degree of control for creating synthetic training data. We express each measure as an independent dimension so that coverage extends to intersectional groups (e.g., “feminine” and “dark skin tone” and “short hair”). When applied to StyleGAN2, QDGS increases the proportion of images recognized with dark skin tones from 9.4% to 25.2%, images recognized as old from 37.0% to 45.7%, and achieves a more uniform spread of gender² presentations and hair length. Notably, QDGS is model-agnostic and adaptable, allowing for its application to various tasks. This is achieved by interchanging the generative model and language prompts according to the desired attributes suited for the training data of the task at hand.

QDGS has the potential to improve trained classifiers. As a proof-of-concept, we show that QDGS can debias an imbalanced dataset of colors and shapes, resulting in up to $\approx 27\%$ improvement on downstream classification accuracy. In the more challenging setting of facial recognition, where models typically perform better for light-skinned users, QDGS is able to increase accuracy on dark-skinned faces from 88.08% to 88.94%. Finally, QDGS achieves the highest average accuracy across facial recognition benchmarks for models trained on biased datasets. These results support training with balanced synthetic datasets.

¹Classifiers trained on Fitzpatrick Skin Type annotations from the IJB-C (Maze et al. 2018) are used to recognize skin tone. Images recognized as types 1 or 2 are denoted as *light*, types 3 or 4 are denoted as *mixed*, and types 5 or 6 are denoted as *dark*.

²We use the term *gender* to refer to a 1D spectrum of facial presentation styles. We denote the ends of the spectrum as *masculine* and *feminine*, and we refer to the center as *androgynous*. Gender attributes, among other continuous demographic attributes, are sometimes discretized into bins for analysis.

2 Ethical Impacts and Limitations

Facial recognition has fraught ethical impacts when misused for the purpose of mass surveillance, and we condemn this use of the technology. Nonetheless, facial recognition technology has been publicly adopted and is used in identity verification systems (e.g., Face ID). Since facial recognition is already integrated in social and institutional contexts, we argue it must be made to be performant for all users.

Facial recognition systems notoriously perform better on light-skinned individuals and worse on dark-skinned individuals (Buolamwini and Gebru 2018). To our knowledge, there are no publicly released facial datasets with skin tone annotations and are balanced across skin tones, which we believe to be the correct approach to assessing the fairness of those systems. We use the Racial Faces in-the-Wild (RFW) dataset (Wang et al. 2019) as a proxy for measuring and addressing this fairness: RFW is approximately balanced across skin tones, but uses the labels Caucasian and African, labeled using the Face++ API. We recognize the flaws in this labeling method and reinterpret these labels as *light-skinned* and *dark-skinned*, respectively.

Debiasing methods often use heuristics that can be useful for some contexts, but not others. While QDGS aims to promote uniform distributions across data attributes, it may not be suited to tasks that necessitate proportionate representations (i.e., when the synthetic dataset’s distribution is similar to that of the target population’s). Further, the language prompts in QDGS enhance the expressivity of desired diversity attributes, but they require careful design to avoid reinforcing linguistic biases. Finally, for contexts that involve vectorized data or motivate human feedback, practitioners may consider using QDGS in conjunction with other methods (Celis et al. 2018; Kou et al. 2021; Ding et al. 2023).

3 Preliminaries

3.1 Learning from Synthetic Data

Advancements in generative models enable the creation of complex and realistic synthetic data, but these improvements often depend on larger models. As generative models scale, their latent spaces also become increasingly complex to navigate, motivating latent space exploration algorithms as part of the sampling process to protect data diversity.

Formally, these latent space exploration algorithms identify inputs $\theta \in \mathbb{R}^n$ for a generative model \mathcal{G} that produce an image $X = \mathcal{G}(\theta)$, optionally maximizing an objective

Dataset	Skin Tone			Age		Gender			Background		Face Rotation		
	Light	Mixed	Dark	Young	Old	Masc	Andro	Fem	Indoors	Outdoors	Frontal	Turned	Profile
Facial Benchmarks (Large Scale)													
LFW	72.7	18.8	08.5	29.6	70.4	75.0	01.0	24.0	60.2	39.9	74.2	23.9	01.9
CelebA	71.5	20.0	08.5	70.9	29.1	41.4	00.5	58.2	49.9	50.1	53.6	41.2	05.2
FFHQ	73.1	17.1	09.8	64.9	35.1	44.2	03.9	51.9	32.0	68.0	55.3	31.2	13.6
IJB-C	59.3	29.7	10.9	33.9	66.1	66.9	01.7	31.3	80.6	19.4	40.5	34.4	25.1
CASIA-WebFace	78.4	09.8	11.9	59.3	40.7	54.9	02.1	42.9	49.8	50.2	71.2	17.2	11.5
Facial Benchmarks (Diversity)													
RFW (Race-diverse)	46.9	25.9	27.2	51.4	48.6	74.8	01.6	23.6	48.7	51.3	69.7	24.3	06.0
BUPT (Race-diverse)	52.6	21.2	26.1	53.4	46.6	67.4	02.6	30.1	46.8	53.2	82.1	12.6	05.3
Synthetic Data Sampling													
Rand50	69.8	20.8	09.4	63.0	37.0	44.0	04.2	51.8	26.2	73.8	54.5	33.0	12.6
QD50 (Ours)	56.8	18.1	25.2	54.3	45.7	51.1	04.3	44.6	26.6	73.4	68.0	26.2	05.8

Table 1: Real datasets contain large imbalances across several demographic and visual attributes. Proportions (%) making up each attribute are listed below. The most uniform distribution of each dataset category is bolded. The proportions are estimated with classifiers trained on IJB-C and CelebA labels, with the computations for each attribute described in Section 6.2.

$f : \mathbb{R}^n \rightarrow \mathbb{R}$. To increase sample diversity, inputs are augmented with additive sample noise: $\theta' = \theta + z$ (Azizi et al. 2023; Tian et al. 2023); similarly, objective functions can factor in diversity by including distance in latent space as a component of the function: $f(\theta') = g(\theta') + \beta \|\theta' - \theta\|_2$ (McDuff et al. 2019; Jain, Memon, and Togelius 2023).

However, large distances in latent space may not correspond to distinct outputs, due to mode collapse in generative models (Bau et al. 2019). Further, latent distances do not directly describe any specific semantic concept, so the resulting variance in data may be ineffective for the downstream task at hand. We refer to distance in latent space as *unmeasured diversity*, since the dimension of variance is ambiguous. While unmeasured diversity is useful for promoting heterogeneity, it does not give practitioners control over the desired concepts for which data should vary. Thus, in addition to unmeasured diversity, we argue that generative sampling algorithms should integrate methods that control *measured diversity* (i.e., dimensions of diversity that are aligned with semantic concepts). In QDGS, we estimate these dimensions with similarity scores to language prompts.

3.2 Fairness in Facial Recognition

Despite potential misuse of facial recognition technology for mass surveillance, its widespread application requires that facial recognition technology be performant for all users. However, large-scale facial datasets tend to disproportionately represent demographic and visual features. Labeled Faces in-the-Wild (LFW) (Huang et al. 2007) contains a $\sim 8.6 : 1$ ratio for light to dark skin tones, a $\sim 3.1 : 1$ ratio for masculine to feminine gender presentations, and a $\sim 2.4 : 1$ ratio for old to young ages (Table 1). Other facial datasets contain similar demographic imbalances, including CelebA (Liu et al. 2015), FFHQ (Karras, Laine, and Aila 2019), IARPA Janus Benchmark-C (IJB-C) (Maze et al. 2018), and CASIA-Webface (CASIA) (Yi et al. 2014).

To improve the transparency of biases in facial data, RFW was designed as a benchmark with separate racial evaluation splits. Subsequently, BUPT-Balancedface (BUPT) (Wang, Zhang, and Deng 2022) was released as a training dataset with racially balanced images. While both of these datasets are estimated to have more uniform distributions of skin tones, they still have high gender presentation disparity with a masculine to feminine ratio of $\sim 3.2 : 1$ in RFW and $\sim 2.2 : 1$ in BUPT. This imbalance reflects the challenge of creating intersectional representations for training datasets.

3.3 Quality-Diversity Optimization

Quality-diversity (Chatzilygeroudis et al. 2021) is a subfield of evolutionary optimization where n -dimensional solutions maximize an objective $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and cover a measure space S formed by the range of k measures $m_j : \mathbb{R}^n \rightarrow \mathbb{R}$, or as a joint measure, $\mathbf{m} : \mathbb{R}^n \rightarrow \mathbb{R}^k$. For each $s \in S$, quality-diversity algorithms aim to find a solution $\theta \in \mathbb{R}^n$ such that $\mathbf{m}(\theta) = s$ and $f(\theta)$ is maximized. In the quality-diversity literature, the task of searching over latent solutions is called latent space illumination (Fontaine et al. 2021).

We employ CMA-MAEGA (Fontaine and Nikolaidis 2023) to solve a relaxed version of quality-diversity by using a tessellation T of measure space S . First, θ is initialized to an initial solution θ_0 . For each of N iterations, a population of λ coefficients c are sampled from a $k + 1$ Gaussian distribution $\mathcal{N}(\mu, \Sigma)$. Next, solutions are branched with objective-measure gradients:

$$\theta' = \theta + |c_0| \nabla f(\theta) + \sum_{j=1}^k c_j \nabla m_j(\theta) \quad (1)$$

An annealed improvement scalar Δ_i is used to rank sampled gradients ∇_i for each new solution θ' based on the update history of a measure cell. We perform gradient ascent on θ weighted on improvement ranking Δ_i . The Gaussian

Algorithm 1: Quality-Diversity Generative Sampling

Input: A generative model \mathcal{G} , objective prompts \mathbf{x}_f , a regularization distance ρ , and measure prompts $\{\mathbf{x}_{m_j}\}_{j=1}^k$.
Parameters: An initial solution θ_0 , the number of iterations N , a learning rate η , a branching population size λ , and an initial branching step size σ_g .
Output: A dataset $\mathcal{D}_{archive}$ with latent solutions for generative model \mathcal{G} .

```

1: Initialize  $\mathcal{D}_{archive} = \phi$ ,  $\theta = \theta_0$ ,  $\mu = \mathbf{0}$ , and  $\Sigma = \sigma_g I$ .
2: for  $iter \leftarrow 1$  to  $N$  do
3:    $X \leftarrow \mathcal{G}(\theta)$ 
4:    $f, \nabla_f \leftarrow \text{compute\_objective}(\theta, X, \mathbf{x}_f, \rho)$ 
5:    $\mathbf{m}, \nabla_{\mathbf{m}} \leftarrow \text{compute\_measures}(X, \{\mathbf{x}_{m_j}\}_{j=1}^k)$ 
6:    $\nabla_f \leftarrow \text{normalize}(\nabla_f)$ ,  $\nabla_{\mathbf{m}} \leftarrow \text{normalize}(\nabla_{\mathbf{m}})$ 
7:    $\mathcal{D}_{archive} \leftarrow \text{update}(\mathcal{D}_{archive}, (\theta, f, \mathbf{m}))$ 
8:   for  $i \leftarrow 1$  to  $\lambda$  do
9:      $\theta'_i \leftarrow \text{branch}(\theta, \nabla_f, \nabla_{\mathbf{m}}, \mu, \Sigma)$ 
10:     $X'_i \leftarrow \mathcal{G}(\theta'_i)$ 
11:     $f', * \leftarrow \text{compute\_objective}(\theta'_i, X'_i, \mathbf{x}_f, \rho)$ 
12:     $\mathbf{m}', * \leftarrow \text{compute\_measures}(X'_i, \{\mathbf{x}_{m_j}\}_{j=1}^k)$ 
13:     $\Delta_i \leftarrow \text{improvement}((f', \mathbf{m}'), \mathcal{D}_{archive})$ 
14:     $\mathcal{D}_{archive} \leftarrow \text{update}(\mathcal{D}_{archive}, (\theta'_i, f', \mathbf{m}'))$ 
15:   end for
16:    $\theta \leftarrow \text{ranked\_ascent}(\eta, \nabla_{1\dots\lambda}, \Delta_{1\dots\lambda})$ 
17:    $\mu, \Sigma \leftarrow \text{adapt}(\Delta_{1\dots\lambda})$ 
18:   if there is no change in  $\mathcal{D}_{archive}$  then
19:     Reset  $\mu = \mathbf{0}, \Sigma = \sigma_g I$ 
20:     Randomly select  $\theta, * \in \mathcal{D}_{archive}$ 
21:   end if
22: end for
23: return  $\mathcal{D}_{archive}$ 

```

parameters μ and Σ are updated with an evolutionary adaptation step. Finally, if no improvement was found, a new seed solution θ is sampled uniformly at random from existing solutions. We visualize these steps in Fig. 2.

4 QDGS with Color-Biased Shapes

To demonstrate QDGS (Algorithm 1), we first use a proof-of-concept colored shapes dataset with a controlled gradation of bias strengths $b \in [0.80, 0.85, 0.90, 0.95, 0.98]$. For each bias strength b , we generate a “real” training dataset such that red triangles and blue squares make up b proportion of the dataset and blue triangles and red squares make up the rest. A variational autoencoder \mathcal{G} is trained on the version with $b = 0.98$. We show that despite the biased generative model \mathcal{G} , QDGS can produce a synthetic dataset that repairs accuracy on a balanced evaluation set with $b = 0.50$.

4.1 Data Generation and Sampling

QDGS finds solutions that maximize an objective function f and measure functions $\{m_j\}_{j=1}^k$.

Objective function: The objective function f computes CLIP similarity scores of generated images with positive and negative language prompts $\mathbf{x}_f = (x_f^{pos}, x_f^{neg})$:

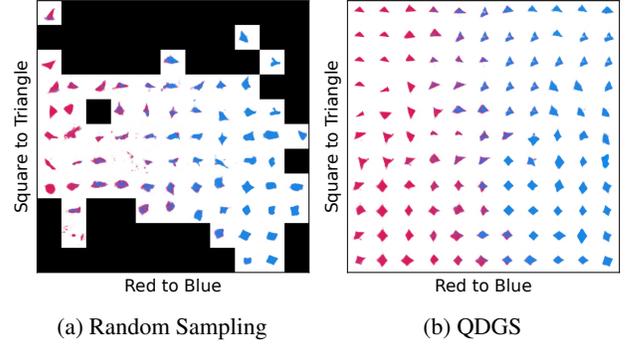


Figure 3: Despite sampling from the same biased red-triangle/blue-square generator, QDGS achieves more uniform coverage over an intersectional measure space.

$$f(\theta) = \text{COS}_{\text{CLIP}}(\mathcal{G}(\theta), x_f^{pos}) - \text{COS}_{\text{CLIP}}(\mathcal{G}(\theta), x_f^{neg}) \quad (2)$$

We use the positive prompt $x_f^{pos} = \text{“A regular square or triangle”}$ and to discourage out-of-distribution images, we use the negative prompt $x_f^{neg} = \text{“Splatters of colors”}$.

Measure functions: Each measure function m_j also computes CLIP scores (Equation 2) to determine the measures of diversity in the resulting images. We run QDGS on $k = 2$ measure dimensions, with the following prompts: $\mathbf{x}_{m_1} = \text{“A \{red, blue\} shape”}$, $\mathbf{x}_{m_2} = \text{“A \{square or diamond with 4 edges, triangle with 3 edges\}”}$. We assign labels via m_2 with a threshold of 0.01.

4.2 Discriminator Training and Evaluation

For each version of our biased dataset, we train a convolutional neural network to discriminate between triangles and squares. Optionally, we fine-tune with the dataset generated with QDGS from the biased autoencoder \mathcal{G} .

4.3 Results for Color-Biased Shapes

Applying QDGS to our proof-of-concept dataset demonstrates two key observations: (1) QDGS is able to generate uniform coverage over intersectional measure spaces; and (2) training with QDGS helps repair biases in classifiers.

In Fig. 3, we show that QDGS is able to generate a balanced proportion of colors with shapes. When randomly sampling, the latent distribution of generator \mathcal{G} spuriously associates red shapes with triangles and blue shapes with squares. Further, generated outputs can sometimes fall out of distribution, resulting in color splatters rather than shapes. Conversely, sampling with QDGS is able to produce balanced images, which are often higher in quality, by evenly distributing red and blue shapes with triangles and squares.

For each bias level, classifier accuracy increases on a fair evaluation benchmark, improving accuracy by up to $\approx 27\%$ (Fig. 1). Notably, QDGS is more helpful when biases in real datasets are stronger. These results suggest that generative models, even when they are biased, are more helpful for training classifiers when used with QDGS.

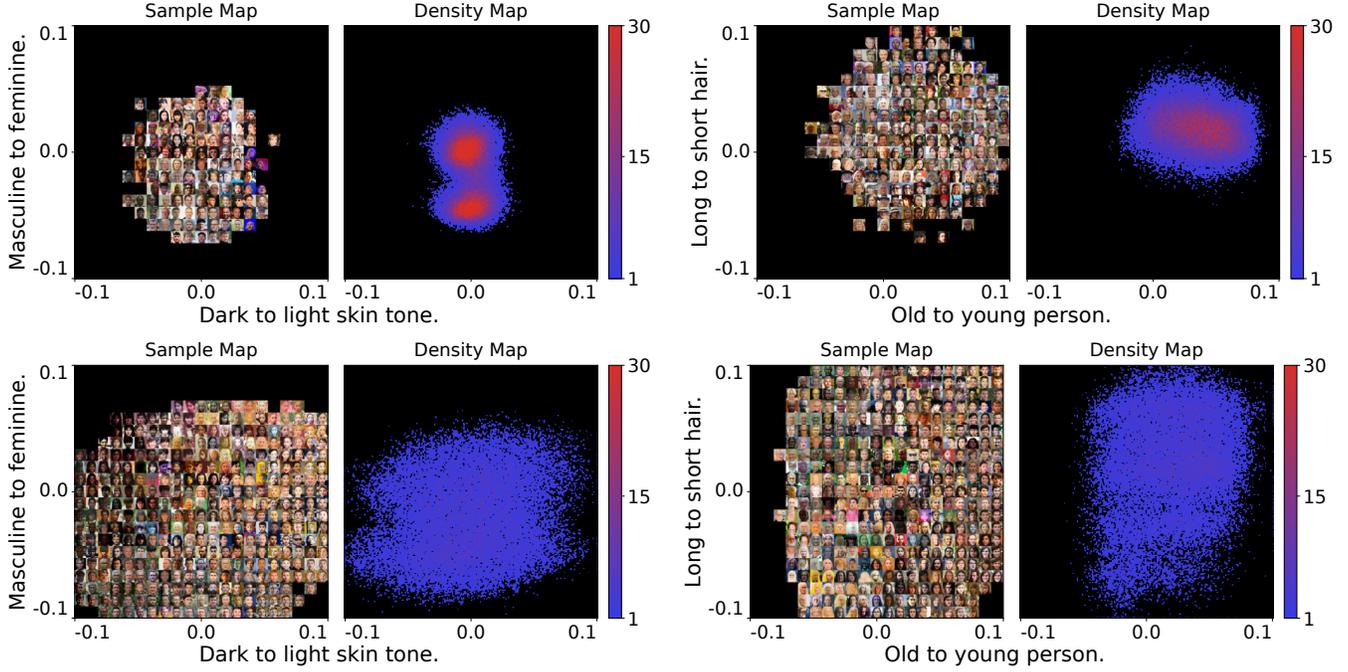


Figure 4: Sample and density maps of synthetic datasets projected onto CLIP latent planes, computed via measure functions described in Section 5.1. Top: Random sampling (Rand50) creates high-density clusters and have sparse representation of extreme measures. Bottom: QDGS (QD50) can create more uniform distributions of synthetic data across measures.

5 QDGS with Facial Recognition

We generate synthetic facial images using the StyleGAN2 model \mathcal{G} with a resolution of 256×256 trained on FFHQ. We create two pretraining datasets with QDGS, QD15 and QD50, from 15K and 50K solutions respectively. We also create randomly sampled baseline datasets Rand15 and Rand50. In all datasets, latent solutions θ are sampled from StyleGAN2’s larger latent space $\mathcal{W} \in \mathbb{R}^{18 \times 512}$.

5.1 Data Generation and Sampling

Objective function: The objective function f is a weighted sum ($\beta_1 = 0.5, \beta_2 = 0.2$) of three component functions, including a text objective g_{txt} , an image margin g_{mgn} , and a regularization factor g_{reg} :

$$f(\theta) = g_{txt}(\mathcal{G}(\theta)) - \beta_1 g_{mgn}(\mathcal{G}(\theta)) - \beta_2 g_{reg}(\theta) \quad (3)$$

The text objective (Equation 2) is used for fine-grained guidance over the semantic quality of visual features. While synthetic image quality can be estimated with other metrics, such as Fréchet inception distance (FID) or a classifier trained on the *Blurry* label of CelebA, these metrics typically do not describe noise structures specific to a domain-task relationship, such as occlusions of the face due to hair or a face mask. Beyond identifying specific noise structures in a domain-task relationship, language prompts can also describe complex, high-level desiderata that would otherwise necessitate the composition of several low-level heuristics. We use the positive prompt $x_f^{pos} = \text{“A detailed photo of an individual with diverse features”}$ and to discourage out-of-

distribution images, we use the negative prompt $x_f^{neg} = \text{“An obscure, fake, or discolored photo of a person”}$.

We incorporate a distance margin to the objective to boost exploration of solutions via unmeasured diversity. We maximize that margin on the subset of randomly selected solutions generated from the past 100 iterations, which we denote $\mathcal{D}_{mem} \subset \mathcal{D}_{archive}$:

$$g_{mgn}(X) = \max_{X' \in \mathcal{D}_{mem}} \cos_{CLIP}(X, X') \quad (4)$$

Finally, we place a constraint on the domain of acceptable solutions by regularizing on the latent space \mathcal{W} . We first identify the mean standardized Euclidean distance δ_{reg} from 10,000 vectors sampled from the latent space \mathcal{W} . We then penalize solutions outside distance δ_{reg} , with a distance multiplier $\rho = 0.5$, where solutions within the boundary distance are not penalized:

$$g_{reg}(\theta) = (\max(\delta(\theta, W), \rho\delta_{reg}) - \rho\delta_{reg})^2 \quad (5)$$

Measure functions: We use Equation 2 to compute the attribute diversity in the resulting images on $k=4$ measure dimensions, with the following prompts: $x_{m_1} = \text{“A \{feminine, masculine\} person”}$, $x_{m_2} = \text{“A person with \{light, dark\} skin”}$, $x_{m_3} = \text{“A person with \{short, long\} hair”}$, and $x_{m_4} = \text{“A person in their \{20, 50\}s”}$.

5.2 Label Assignment and Data Augmentation

To assign labels while maintaining feasibility for memory, we first separate each dataset into equally-sized 3×3 chunks based on the first two measures, m_1 and m_2 . Within each

chunk $\mathcal{C}_j \subset \mathcal{D}_{archive}$, we perform K-means clustering on latent solution θ where $k = \frac{|\mathcal{C}_j|}{2}$, and each cluster is assigned a separate identity label. For each latent solution θ , we perform latent walks along directions that correspond to horizontal facial rotation and affective valence, creating 9 total variants of each image.

5.3 Discriminator Training and Evaluation

We select ResNet101 backbones and train with AdaFace loss (Kim, Jain, and Liu 2022). For each synthetic dataset, we pretrain model weights over 26 epochs with a learning rate of 0.01. Each model is trained on the skin-tone imbalanced CASIA and skin-tone balanced BUPT datasets. We run 10 trials for each training configuration.

To analyze facial recognition performance on skin-tone-split evaluation sets, we first evaluate on RFW. Next, we evaluate on LFW, CFPFP (Sengupta et al. 2016), CPLFW (Zheng and Deng 2018), CALFW (Zheng, Deng, and Hu 2017), and AgeDB (Moschoglou et al. 2017). We follow the testing protocols of each respective benchmark.

6 Results and Discussion

6.1 Spread of Synthetic Data Distributions

QDGS is able to produce a more uniform spread of synthetic data compared to random sampling from StyleGAN2. Randomly sampled image distributions of StyleGAN2 consist of high density clusters with one or more modes, but these modes may arise from sampling biases rather than reflecting genuine modes of real populations. For example, while real skin tone distributions have multiple modes, randomly sampled images from StyleGAN2 tend to collapse into a single mode with respect to skin tone measures (Fig. 4). Conversely, QD50 exhibits higher variance and covers a broader area of the measure space compared to Rand50. These observations are vital when training with synthetic data, as models are susceptible to overfitting on dominant classes or features.

Sampling with QDGS can also inform relationships between generative latent spaces and semantic concepts. The marginal distributions indicate that randomly sampling $\theta \in \mathcal{W}$ naturally captures a wide range of genders and ages, but underrepresented skin tones (very “light” and very “dark”) and hair lengths (“long hair”) are considered outliers of the distribution of \mathcal{W} . Further, interactions between measures can reveal associations represented in StyleGAN2 and CLIP. For example, image data with high similarity to language prompts (“dark skin” and “feminine”), (“light skin” and “masculine”), and (“in their 20s” and “long hair”) are sparse in QD50. These spurious associations in synthetic datasets can amplify stereotypes in downstream tasks.

Visual inspection of the image sample maps show that images in QD50 have more extreme visual features than those in Rand50. While these data contain more varied and possibly out-of-distribution attributes, this unmeasured diversity can help facilitate robustness and domain adaptation to unseen test cases. For example, a model that is trained on images with neon lighting (QD50) can learn important information on lighting and colors that may help it extrapolate to predicting test cases that are in grayscale (AgeDB).

Pretraining	Dark-skinned	Light-skinned	DI
Train: CASIA (Imbalanced)			
None	88.08 ± 0.07	94.05 ± 0.06	93.65 ± 0.10
Rand15	88.17 ± 0.08	94.41 ± 0.07	93.38 ± 0.14
Rand50	88.69 ± 0.10	94.67 ± 0.07	93.68 ± 0.10
QD15 (Ours)	88.25 ± 0.09	94.42 ± 0.06	93.47 ± 0.10
QD50 (Ours)	88.94 ± 0.07	94.62 ± 0.10	93.99 ± 0.06
Train: BUPT (Balanced)			
None	96.22 ± 0.03	97.58 ± 0.05	98.60 ± 0.07
Rand15	96.22 ± 0.05	97.72 ± 0.05	98.47 ± 0.09
Rand50	96.17 ± 0.09	97.82 ± 0.04	98.31 ± 0.08
QD15 (Ours)	96.29 ± 0.06	97.75 ± 0.03	98.51 ± 0.07
QD50 (Ours)	96.32 ± 0.06	97.80 ± 0.04	98.49 ± 0.06

Table 2: Pretraining with QDGS can improve accuracy on the dark-skinned subgroup in RFW. Values represent the averaged accuracy (%) and disparate impact ratio (%) with standard error metrics, aggregated for 10 trials.

Tone	Baseline	QDGS	AdaFace	QDGS + AdaFace
Dark	79.7 ± 0.3	82.8 ± 0.2	88.1 ± 0.1	88.9 ± 0.1
Light	88.8 ± 0.1	90.8 ± 0.3	94.1 ± 0.1	94.6 ± 0.1

Table 3: Ablation results: QDGS is helpful for either loss.

6.2 Generalization of Language Guidance to Categorical Diversity Labels

In addition to maximizing language-guided diversity, QDGS generalizes to unseen categorical diversity labels. We classify data using models trained on IJB-C (*Skin Tone, Background, Face Rotation*) and CelebA (*Age, Gender*), where each prediction probability is discretized into two (*Background, Age*) or three (*Skin Tone, Gender, Face Rotation*) bins. In Table 1, we present these proportions for facial benchmarks and our synthetic datasets Rand50 and QD50.

For labels corresponding to measures (e.g., *Skin Tone* : m_1), QD50 is able to achieve a more uniform distribution than Rand50. QDGS increases the proportion of images recognized as dark skin tones from 9.4% to 25.2%. QDGS is also able to simultaneously increase the proportion of images recognized as old from 37.0% to 45.7%. These results demonstrate that balancing effects from QDGS generalizes to predictive distributions for CelebA and IJB-C labels (e.g., *Fitzpatrick Skin Type 6*), even though QDGS uses language prompted measures (e.g., “dark skin”). Generalization to categorical diversity labels supports that language prompting can be valuable approximations of diversity without the use of additional labeled data.

However, it is intractable to anticipate an exhaustive list of the desired diversity measures *a priori*. This challenge is demonstrated by facial rotation, which has a high skew in QD50. QDGS also relies on a generative model’s ability to generate desired concepts, which may not be the case when the data are simply not represented to a sufficient standard. In these scenarios, QDGS will be unsuccessful at generalizing to the corresponding categorical labels.

Pretraining	LFW	CFPP	CPLFW	CALFW	AgeDB	AVG
Train: CASIA (Imbalanced)						
None	99.38 ± 0.02	95.14 ± 0.06	90.24 ± 0.09	93.53 ± 0.05	94.33 ± 0.06	94.52 ± 0.03
Rand15	99.48 ± 0.02	95.68 ± 0.05	90.69 ± 0.07	93.62 ± 0.03	94.67 ± 0.09	94.83 ± 0.03
Rand50	99.45 ± 0.02	95.70 ± 0.04	90.91 ± 0.08	93.60 ± 0.04	94.82 ± 0.05	94.89 ± 0.03
QD15 (Ours)	99.45 ± 0.02	95.67 ± 0.04	90.78 ± 0.06	93.56 ± 0.04	94.67 ± 0.06	94.83 ± 0.03
QD50 (Ours)	99.50 ± 0.01	95.72 ± 0.03	90.94 ± 0.06	93.71 ± 0.06	94.72 ± 0.07	94.92 ± 0.02
Train: BUPT (Balanced)						
None	99.71 ± 0.02	96.15 ± 0.04	93.00 ± 0.04	95.60 ± 0.02	96.87 ± 0.05	96.27 ± 0.02
Rand15	99.74 ± 0.02	96.29 ± 0.05	93.17 ± 0.05	95.67 ± 0.04	96.90 ± 0.02	96.35 ± 0.02
Rand50	99.75 ± 0.02	96.17 ± 0.06	93.07 ± 0.04	95.65 ± 0.04	96.93 ± 0.05	96.31 ± 0.03
QD15 (Ours)	99.75 ± 0.01	96.23 ± 0.04	93.04 ± 0.07	95.66 ± 0.03	96.93 ± 0.05	96.32 ± 0.03
QD50 (Ours)	99.73 ± 0.02	96.21 ± 0.05	93.08 ± 0.06	95.59 ± 0.03	96.90 ± 0.07	96.30 ± 0.02

Table 4: Pretraining with QDGS produces potential improvements in accuracy. Values represent averaged accuracy (%) with standard error metrics evaluated on five evaluation benchmarks, aggregated for 10 trials.

6.3 Repairing Biases in Classifiers with QDGS

Using QDGS, we show that facial recognition models can be more performant on faces with dark skin tones. We compute the accuracy on reinterpreted RFW test splits and the disparate impact ratio (DI) for dark-skinned to light-skinned data (Table 2). DI is represented as the ratio of performance on a minority group to that of a majority group, where values closest to 100% indicates equal treatment of groups.

Models pretrained with QD15/50 achieve higher accuracies for dark-skinned faces and the same accuracies for light-skinned faces compared to those pretrained with Rand15/50. On models trained on skin-tone imbalanced data (CASIA), QDGS achieves the highest accuracy for dark-skinned faces, raising scores from 88.08% to 88.94%. On models trained on skin-tone balanced data (BUPT), pretraining with QDGS has a smaller but still noticeable improvement on dark-skinned faces, achieving the highest accuracy for dark-skinned faces across training methods. These trends are reflected in DI, with QD15/50 consistently outperforming Rand15/50, especially when models are trained on imbalanced data. These results indicate that QDGS is helpful for pretraining, particularly when training data are imbalanced.

AdaFace, an adaptive loss that assigns higher weight to difficult training examples, may have some debiasing properties. To explore whether QDGS is effective without using an adaptive loss, we conduct an ablation (Table 3). For our baseline, we run on the ArcFace loss (Deng et al. 2019). In both cases, using QDGS improves performance on dark-skinned faces (79.7% to 82.8% on ArcFace and 88.1% to 88.9% on AdaFace), suggesting that QDGS is helpful when used in conjunction with other debiasing methods.

6.4 Effects of QDGS on Accuracy

QDGS can maintain or improve accuracy, and we do not observe accuracy decline on any evaluation benchmark. Since all pretrained models perform better than no pretraining, we believe that synthetic data can help to prime models with diverse data attributes. Thus, the models are equipped to generalize across distributional discrepancies when encountering novel data, surpassing models without pretraining.

On models trained with CASIA, QDGS achieves a higher average accuracy than randomly sampled pretraining (Table 4). These differences are larger between QD50 and Rand50 than for QD15 and Rand15. On models trained with BUPT, there are fewer differences in accuracy based on the pretraining method. These trends likely reflect the larger scale of BUPT ($\approx 1.3M$ images in BUPT v.s. $\approx 200K$ images in CASIA), since training on a larger dataset can diminish the effects of pretraining. Another possibility is that BUPT contains more heterogeneous data and therefore experiences fewer added benefits from QDGS.

7 Conclusions

This work presents QDGS, a framework for sampling balanced synthetic training datasets from generative models. QDGS can increase spread over desired attributes, described by user-defined language prompts, by exploring high dimensionality latent spaces. With QDGS, we are able to repair biases in shape classifiers up to $\approx 27\%$. In the more challenging domain of facial recognition, we are able to achieve a more uniform spread of synthetic facial image training data, simultaneously increasing the proportion of images recognized with dark skin tones from 9.4% to 25.2% and images recognized as old from 36.0% to 45.7%. For models trained on biased datasets, QDGS improves accuracy performance on dark-skinned faces from 88.08% to 88.94%, and QDGS achieves the highest average accuracy across facial recognition benchmarks. These results support training with balanced synthetic datasets.

Acknowledgements

This product contains or makes use of the following data made available by the Intelligence Advanced Research Projects Activity (IARPA): IARPA Janus Benchmark C (IJB-C) data detailed at Face Challenges homepage. The authors acknowledge the Center for Advanced Research Computing (CARC) at the University of Southern California for providing computing resources that have contributed to the research results reported within this publication. URL: <https://carc.usc.edu>.

References

- Azizi, S.; Kornblith, S.; Saharia, C.; Norouzi, M.; and Fleet, D. J. 2023. Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*.
- Balestriero, R.; Bottou, L.; and LeCun, Y. 2022. The effects of regularization and data augmentation are class dependent. *Neural Information Processing Systems (NeurIPS)*.
- Bau, D.; Zhu, J.-Y.; Wulff, J.; Peebles, W.; Strobel, H.; Zhou, B.; and Torralba, A. 2019. Seeing what a gan cannot generate. In *International Conference on Computer Vision (ICCV)*.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Fairness, Accountability and Transparency (FAccT)*.
- Celis, E.; Keswani, V.; Straszak, D.; Deshpande, A.; Kathuria, T.; and Vishnoi, N. 2018. Fair and diverse DPP-based data summarization. In *International Conference on Machine Learning (ICML)*.
- Chatzilygeroudis, K.; Cully, A.; Vassiliades, V.; and Mouret, J.-B. 2021. Quality-Diversity Optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*. Springer.
- Cho, J.; Zala, A.; and Bansal, M. 2023. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. In *International Conference on Computer Vision (ICCV)*.
- Deng, J.; Guo, J.; Xue, N.; and Zafeiriou, S. 2019. Arcface: Additive angular margin loss for deep face recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
- Ding, L.; Zhang, J.; Clune, J.; Spector, L.; and Lehman, J. 2023. Quality Diversity through Human Feedback. *arXiv preprint arXiv:2310.12103*.
- Fontaine, M.; and Nikolaidis, S. 2023. Covariance matrix adaptation map-annealing. In *Genetic and Evolutionary Computation Conference (GECCO)*.
- Fontaine, M. C.; Liu, R.; Khalifa, A.; Modi, J.; Togelius, J.; Hoover, A. K.; and Nikolaidis, S. 2021. Illuminating mario scenes in the latent space of a generative adversarial network. In *AAAI Conference on Artificial Intelligence*.
- Friedrich, F.; Schramowski, P.; Brack, M.; Struppek, L.; Hintersdorf, D.; Luccioni, S.; and Kersting, K. 2023. Fair diffusion: Instructing text-to-image generation models on fairness. *arXiv preprint arXiv:2302.10893*.
- Huang, G. B.; Ramesh, M.; Berg, T.; and Learned-Miller, E. 2007. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. Technical report, University of Massachusetts, Amherst.
- Jain, A.; Memon, N.; and Togelius, J. 2023. Zero-shot racially balanced dataset generation using an existing biased StyleGAN2. *arXiv preprint arXiv:2305.07710*.
- Jain, N.; Olmo, A.; Sengupta, S.; Manikonda, L.; and Kambhampati, S. 2022. Imperfect ImaGANation: Implications of GANs exacerbating biases on facial data augmentation and snapchat face lenses. *Artificial Intelligence*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Computer Vision and Pattern Recognition (CVPR)*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of stylegan. In *Computer Vision and Pattern Recognition (CVPR)*.
- Kim, M.; Jain, A. K.; and Liu, X. 2022. Adaface: Quality adaptive margin for face recognition. In *Computer Vision and Pattern Recognition (CVPR)*.
- Kou, Z.; Zhang, Y.; Shang, L.; and Wang, D. 2021. Faircrowd: Fair human face dataset sampling via batch-level crowdsourcing bias inference. In *International Workshop on Quality of Service (IWQoS)*.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *International Conference on Computer Vision (ICCV)*.
- Maze, B.; Adams, J.; Duncan, J. A.; Kalka, N.; Miller, T.; Otto, C.; Jain, A. K.; Niggel, W. T.; Anderson, J.; Cheney, J.; et al. 2018. Iarpa janus benchmark-c: Face dataset and protocol. In *International Conference on Biometrics (ICB)*.
- McDuff, D.; Ma, S.; Song, Y.; and Kapoor, A. 2019. Characterizing bias in classifiers using generative models. *Neural Information Processing Systems (NeurIPS)*.
- Moschoglou, S.; Papaioannou, A.; Sagonas, C.; Deng, J.; Kotsia, I.; and Zafeiriou, S. 2017. Agedb: the first manually collected, in-the-wild age database. In *Faces “In-The-Wild” Workshop-Challenge, Computer Vision and Pattern Recognition (CVPR)*.
- Otterbacher, J. 2018. Social cues, social biases: stereotypes in annotations on people images. In *Human Computation and Crowdsourcing (HCOMP)*.
- Sengupta, S.; Chen, J.-C.; Castillo, C.; Patel, V. M.; Chellappa, R.; and Jacobs, D. W. 2016. Frontal to profile face verification in the wild. In *Winter Conference on Applications of Computer Vision (WACV)*.
- Tian, Y.; Fan, L.; Isola, P.; Chang, H.; and Krishnan, D. 2023. StableRep: Synthetic Images from Text-to-Image Models Make Strong Visual Representation Learners. *arXiv preprint arXiv:2306.00984*.
- Wang, M.; Deng, W.; Hu, J.; Tao, X.; and Huang, Y. 2019. Racial Faces in the Wild: Reducing Racial Bias by Information Maximization Adaptation Network. In *International Conference on Computer Vision (ICCV)*.
- Wang, M.; Zhang, Y.; and Deng, W. 2022. Meta Balanced Network for Fair Face Recognition. *Transactions on Pattern Analysis and Machine Intelligence*.
- Yi, D.; Lei, Z.; Liao, S.; and Li, S. Z. 2014. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*.
- Zheng, T.; and Deng, W. 2018. Cross-pose LFW: A database for studying cross-pose face recognition in unconstrained environments. Technical Report 18-01, Beijing University of Posts and Telecommunications.
- Zheng, T.; Deng, W.; and Hu, J. 2017. Cross-age lfw: A database for studying cross-age face recognition in unconstrained environments. *arXiv preprint arXiv:1708.08197*.